XC System Software Version 2.1 is an HP supported software stack which can be used for the operation and management of an HP Cluster Platform. Hereafter in this document, the term *XC System Software* will refer to XC System Software V2.1, and the combination of HP Cluster Platform with XC System Software will be called an XC System. The XC System is based on HP Integrity or HP ProLiant nodes, connected by a high-speed interconnect (XC System Interconnect) such as the Quadrics QsNetII switch network from Quadrics Ltd, the Myrinet-2000 switch network from Myricom, an Infiniband switch network, or a Gigabit Ethernet switch network. In addition to the XC System Interconnect, each XC System also includes an administrative network (XC Administrative Network) based on Gigabit and 10/100 Mbit Ethernet switches with separate console and, optionally, administrative connections to each node. The XC System Software includes HP MPI, HP MLIB and other HP proprietary and open source software that supplies the infrastructure for optimal interconnect performance. In addition, XC System Software includes a Linux distribution compatible with Red Hat Enterprise Linux Advanced Server 3.0, and software that provides users and system administrators with single-system attributes for managing and using XC system resources. Also, XC System Software includes technology from Platform Computing Inc. (Platform) for job management and optimized throughput of parallel applications.

Product Overview

XC System Software creates a performance optimized application environment that is manageable across a range of system scales. XC System Software is based on a standard Linux distribution combined with a number of open systems packages that leverages the work of the overall open systems community. This open systems base combined with technology from Hewlett-Packard and its partners, achieves a powerful and complete solution for development, execution and management of simultaneous parallel and serial applications. Preconfigured systems support up to 512 nodes; larger systems may require custom configuration. The XC System Software is an integration of premier components from leading suppliers of cluster technologies. In selecting these components, several key criteria have been employed:

- Capability - the components are capable of providing, in total, the features required for a complete high-performance cluster solution
- Provability - the components have been widely used and proven in production facilities
- Integrability - the components are capable of being integrated into a single management and deployment environment with uniform semantics, and without multiple conflicting methods for job management or execution
- Portability - the components can be combined to produce a whole system that can assure application portability across the platform family, and can adopt to future technologies

Typical usage is described in what follows. The XC System can be integrated into an existing network environment and presented to users, including system administrators, as a single hostname and IP address on the external network. When a user logs into the XC system, the login session is directed to one of several nodes which have been enabled to accept login sessions. Upon logging into the system, the users may then develop, execute and manage their applications. For development, the environment consists of standard Linux tools such as gcc and gdb and other third-party, as well as optimized parallel libraries that are part of the XC System Software. These libraries include the HP MLIB math library and the HP MPI message-passing library. HP MPI provides binary compatibility across interconnects. The implementations of HP MPI and HP MLIB may invoke low-level interconnect-specific libraries such as the Quadrics QSNet or the Myricom GM libraries, but entirely isolate the user from the requirement to directly access those libraries, or even to be aware of them.

All nodes may be used to run user applications but one or more of the nodes may also be used for other purposes such as login services, job launch, and system administration. Each node runs a full Linux environment. For job execution and policy centric workload management, XC System Software provides an integrated resource management, scheduling and job launch mechanism based on the underlying technologies LSF® HPC software ("LSF") and *SLURM*. *SLURM* is responsible for the allocation of resources to jobs. The resources are provided by nodes that are designated as compute resources. Processes of the job are assigned to (executed on) those allocated resources. As part of that process assignment / setup, the original submission environment is replicated within those allocated resources. *SLURM* also provides the means to arrange an XC System into distinct partitions. Serial or parallel jobs may be scheduled for execution within a given partition, provided that the partition has sufficient resources (for example, memory, or number of CPUs (a.k.a. cores)) to execute that job. The entire system can be designated as a single partition, allowing parallel jobs to run across all of the CPUs of the XC System. Alternatively, the system administrator can divide the system into smaller partitions, or can set up the system so that users bypass *SLURM* altogether. If the system is configured to bypass *SLURM*, then workload management can be accomplished by standard LSF (available from Platform Computing Inc.) which is also licensed as part of XC

Overview

System Software.

Application file I/O can be performed to locally mounted storage devices, or alternatively, to remote storage using NFS or *Lustre*- for high performance and high availability - to access the remote file systems. Intra-XC NFS operations can be configured to use either the XC Administrative Network or the XC System Interconnect. By using a separate interconnect for administration and I/O, the XC System administrator is able to isolate user application traffic from administrative operations and monitoring. With this separation, application I/O performance and process communication can be made more predictable while still enabling administrative operations to proceed. Alternatively, the XC System administrator may prefer to route I/O across the administrative network so that the I/O has less impact on the synchronicity of parallel applications. The XC System Software provides the flexibility to support such administrative needs. Any XC System node is permitted to export local files to NFS clients connected to the XC Administrative Network. If the node is also connected to an external network, local files can be exported to an external NSF client. In order to reduce the impact of I/O on parallel computations, administrators may choose to organize the system so that I/O nodes are separately partitioned from compute nodes.

System administrators perform the majority of their administrative tasks from a node designated as the XC head node (head node) and which is used to initiate booting the XC System through the process of XC cluster discovery. This head node is also capable of providing system services through hosting daemons such as the *SLURM* and LSF masters. It is also possible to configure the XC System so that system services are distributed amongst several nodes, rather than centralized in a single head node. This distribution of services might be advised for reasons of system availability, scalability, heavy system load, and so on.

The XC System Software installation process is designed to provide the maximum amount of automation with the minimal amount of user intervention. Software installation is done by first loading software onto the XC head node and then distributing the system software to the other nodes using the SystemImager utility. This distribution is performed over the XC Administrative Network. XC configuration utility software is used during the initial configuration, to customize certain system files on each node, based on data gathered during the cluster discovery phase, and stored in the XC Configuration and Management Database. Subsequent to initial configuration, the administrator can perform and log console commands such as starting, stopping, and booting nodes from the head node using the XC Console Management Facility (CMF) and communicating over the console branch of the XC Administrative Network. (CMF is available only on Xeon and Itanium-based systems. For Opteron-based systems, the administrator performs console operations using telnet. Console-logging is not available for the Opteron-based systems.) The administrator uses a single user-interface to report statistics and to post events based on data collected by scalable software that collects usage and health information (hardware parameters) from the individual nodes. The XC software is delivered with pre-configured rules to monitor all critical systems and to report the information through one centralized interface. The monitoring technology for data-collection is *SuperMon* and the user-interface technology is *Nagios*. *Nagios* events can be monitored with a browser, as well as a command-line interface. XC System Software also provides a scalable event-logging utility, *syslog-ng*. This utility aggregates cluster events of interest to a single location. In addition to monitoring and logging global information, it is occasionally necessary for the administrator to perform actions on every node, or on some subset of nodes, of the system. This capability is supplied by the global parallel shell utility, *pdsh*.

Technology Components

XC System Software integrates technology that has been widely tested and used by high performance computing sites. The key technology components of XC System Software are based on:

- Open Source Technology Components (see the table and the later section on Linux)
- LSF Version 6.1 and LSF HPC Version 6.1
- HP MPI Version 2.1.1
- HP MLIB Version 1.5
- Console Management Facility (from HP AlphaServer System Software V2.6)

The following table lists the key open source components that have been integrated together to provide the clustering and parallel computing features of the XC System Software. The table shows the version number and source URL used for integration in the XC System Software. In addition to these components is the Linux distribution itself, including the packages (RPMs) listed later. All of this technology is covered by HP software support contracts for XC System Software, provided that the only changes made to the software are those authorized by Hewlett-Packard.

Key Open Source Technology Components		
Name	Version	URL
MySQL	4.0	http://www.mysql.com
SystemImager	3.2.3	http://www.systemimager.org
Linux Virtual Server	1.21-9	http://www.linuxvirtualserver.org
SLURM	0.3.8-1	http://www.llnl.gov/linux/slurm
Nagios®	1.2	http://www.nagios.org
Nagios Plugins	1.3.1	http://www.nagios.org
Nrpe	2.0	http://www.nagios.org
Nsca	2.4	http://www.nagios.org
SuperMon	1.4.6	http://supermon.sourceforge.net
Syslog_ng	1.6.2.1	http://www.balabit.com/products/syslog_ng
Parallel Distributed Shell	1.7-6	http://www.llnl.gov/linux/pdsh/pdsh.html
Modules	3.1.6	http://modules.sourceforge.net

Installation and Configuration

XC System Software is delivered on a single DVD, which includes a default *kickstart* file to enable the installation to proceed automatically with Red Hat style installation. The DVD is placed in the node that will be the head node. The administrator starts up the Linux installation, and when prompted, selects from a list, a hard drive to be used for installation. The administrator also needs to provide some networking information and the root password. The installation then proceeds to completion on the head node, without any further input from the administrator, and without requiring swapping of DVD's or CD's. If the user requires additional packages (RPMs), these can be installed before the cluster configuration is performed. After installing software on the head node, the configuration phase begins through the process of automated cluster topology discovery, a process facilitated through the pre-defined wiring schemes for the XC Systems. Information gathered during the discovery phase includes network topology, console addresses, network addresses and the binding between the XC System Interconnect and XC Administrative networks. Configuration continues through a simple user interface, with information gathered on the roles to be played by various nodes, such as NIS data when applicable, and so on. The configuration data is written to the XC Configuration and Management Database (CMDB), for subsequent retrieval by various utilities. The utilities access the CMDB via SQL commands, as provided by MySQL.

In the next step of the installation, a collection of head-node system files is assembled for installation on the other nodes. This assemblage, known as the Base Image, is accomplished with SystemImager. The nodes are each booted over the network (DHCP boot) with an autoinstall kernel which sets up local file systems. The Base Image is then distributed to all the nodes. In the final phase of configuration, node-specific configuration information is imparted to each node, based in part on data in the CMDB. Each node ultimately boots and runs its copy of the operating system, using system files on a local disk. Once the system has been configured,

global updates can be efficiently performed through XC System Software utilities which modify the Base Image and which then distribute the changed files to all of the nodes or to selected nodes. Node-specific updates are done through utilities that modify the CMDB, and then accessing this data the next time the affected nodes are rebooted.

Single System Access

The *Linux Virtual Server* (LVS) is used to present a single host name to users so that they see the XC System as a single system for login access to prepare and launch high-performance applications. LVS directs each user login session to one of several potential login nodes. LVS can be used on nodes which have external network connections in additions to the connections that are mandatory for the XC. The XC System architecture does not restrict the number or types of nodes that have such external connections, nor does the architecture mandate the use of LVS for such nodes. Nodes without external network connections cannot be opened from external addresses, but the XC System Software optionally supports Network Address Translation (NAT) which allows all nodes in the system to open network connections to external addresses. NAT can be used, for example, to include the XC System as a member of a NIS (Network Information Services) domain.

File System

Each node of the XC System has its own local copy of all the XC System Software files including the Linux distribution and also has its own local user files. Every node may also import files from NFS or *Lustre* fileservers. XC System Software supports NFS 3 including both client and server functionality. XC System Software also enables *Lustre* client services for high-performance and high-availability file I/O. These *Lustre* client services require the separate installation of *Lustre* software provided with the HP StorageWorks Scalable File Share (SFS).

In the case of NFS files, these can be shared purely between the nodes of the XC System, or alternatively can be shared between the XC System and external systems. External NFS files can be shared with any node having a direct external network connection. It is also possible to set up NFS to import external files to XC nodes without external network connections, by routing through a node with an external network connection. The administrator can choose to use either the XC Administrative Network, or the XC System Interconnect, for NFS operations. The XC System Interconnect can potentially offer higher performance, but only at the potential expense of the performance of application communications.

For high-performance and/or high-availability file I/O, the *Lustre* file system should be used. The *Lustre* file system uses POSIX-compliant syntax and semantics. The XC System Software includes kernel modifications required for *Lustre* client services which enable the operation of the separately installable *Lustre* client software. The industry-leading external *Lustre* file server product is the HP StorageWorks SFS, which fully supports the XC System. The HP Scalable File Share (SFS) V2.0-0 includes XC *Lustre* client software. The SFS can be integrated with the XC System so that *Lustre* I/O is performed over the same high-speed interconnect fabric used by the XC system. So, for example, if the XC System Interconnect is based on a Quadrics QsNet^{II} switch, then the SFS will serve files over ports on that switch. The file operations are able to proceed at the full bandwidth of the XC System Interconnect since these operations are implemented directly over the low-level communications libraries. Further optimizations of file I/O can be achieved at the application level using special file system commands - implemented as iocts - which allow a program to interrogate the attributes of the file system, modify the stripe size and other attributes of new (zero-length) files, and so on. Some of these optimizations are implicit in the HP MPI I/O library, which implements the MPI-2 file I/O standard.

Workload and Resource Management

XC System Software provides both direct job allocation and policy-scheduled interactive and batch queue management based on Platform's software product - LSF HPC Version 6.1. Alternatively, an XC System can be configured to bypass the use of *SLURM* and to use standard LSF Version 6.1. Further information regarding this alternative can be found at http://docs.hp.com/en/linux.html. When using standard LSF, all features of that product (see, for example, www.platform.com) are fully supported on the XC System The remainder of this section will describe the combined use of *SLURM* and LSF HPC Version 6.1(LSF). The expected usage of the XC System is through scheduled, rather than direct, access. LSF provides a rich set of scheduling policies for scheduling and prioritizing jobs based on combinations of static and dynamic system attributes, such as the number of processors, job attributes, such as time limits, and user-attributes, including uid, gid, account id, project id, and priority-based quotas for size and time allocation. LSF policies include first-come-first-served, fairshare, hierarchical fairshare, deadline-constrained, exclusive, adaptive dispatch, backfill, preemptive, and job slot reservation. An extremely reliable, and richly configurable network-based queuing system provides maximum flexibility and central administration of all work in the cluster while ensuring that all work submitted to the XC System runs to completion. LSF provides comprehensive job management such as tracking and control. An extensible LSF checkpoint interface supports application-level checkpointing on the XC System, when supported by the application. In addition, LSF tracks job ids and elapsed time on the XC System, and provides an accounting log for system utilization and chargeback analysis. Accounting information includes submit, start and stop time (waiting time in the queue can be calculated), total user time, total system time, total maximum memory usage (physical and virtual), job id, user id, project id and number of CPUs used by the job.

The *SLURM* resource manager allocates access to nodes within the XC System, launches parallel and serial jobs, manages those jobs, and provides I/O control, signals, limits and so on. *SLURM* allows the administrator to treat the total set of CPUs, provided by all nodes, as a single large pool. Alternatively, this set of nodes can be partitioned into a partition for use by LSF jobs, and/or one or more other partitions, where users can directly submit jobs without the use of LSF. *SLURM* is responsible for finding, allocating and utilizing a set of nodes that fits the job's requirements. *SLURM* supports node 'features' predefined by the system administrator that can be specifically requested by user jobs during submission. These features are useful for identifying compute nodes with special characteristics, such as more memory, faster CPUs, additional local disk space, or any additional hardware installed on those nodes. These *SLURM* 'features' are also supported, with XC System Software, by LSF.*SLURM* also creates the requisite processes on each node. Non-LSF partitions would typically be set up for different types of jobs such as development work, or to set aside nodes with different types of characteristics for special-purpose use. Jobs cannot span across nodes assigned to different partitions. The XC software provides the ability to prevent users from logging into nodes that they have not previously allocated through *SLURM*. This ensures exclusive access to nodes for parallel applications that could otherwise be affected by any additional undesired activity from other users.

Scalable application resource management is implemented through the combination of a control daemon (slurmctld) responsible for the entire XC System, and a collection of client daemons (slurmd) running on each of the XC nodes. The *SLURM* control daemon manages the queue of pending jobs, provides node state information, and allocates nodes. The client daemons are responsible for providing node status, job status, implementing remote execution, stream copy (stdin, stdout and stderr), and job control (signals). *SLURM* provides capabilities for access control, and checks resource limits before scheduling the user's jobs. The *SLURM* srun command is used, when submitting jobs to the direct-launch partition, to allocate resources, initiate jobs that run interactively, execute jobs submitted by LSF, or attach to running jobs. Options to srun include the processor count, the node count, specific nodes to use or avoid, nodes that can be shared and so on. *SLURM* also provides an interface (sinfo) for viewing the state of partitions or nodes - for example, which nodes are currently available. If the user wishes to terminate a job or send an arbitrary signal, that is done with the *SLURM* scancel command, which can specify a user, a program name, a partition, job state, etc. The *SLURM* squeue command is used to monitor existing *SLURM* jobs.

LSF is integrated with *SLURM* to combine the use of *SLURM*'s extensive launching options and resource knowledge with LSF's powerful scheduling to create a high-quality, comprehensive workload and resource management system on XC. The LSF scheduler obtains all node-usage information from *SLURM*, and then uses that information to submit appropriately-scheduled jobs to *SLURM*.

Standard Features

The following LSF Version 6.1 features are not supported in XC System Software:

- Interactive remote execution (LSF Base commands) and load-sharing shell
- The following LSF job limits are not supported: swap, process, and thread
- Distinct runtime resource utilization such as cpu time and I/O time (although users can launch their own resource monitoring programs if necessary)
- Per-node load indice support (given the CPU-based scheduling on XC, CPUs are expected to be either fully loaded with a user's application or idle waiting for a job)
- LSF Analyzer

Administration

The XC System Software administrative tools provide mechanisms for configuring (see above), monitoring, controlling, and managing a node or groups of nodes. Nodes are monitored using *Nagios*. *Nagios* is a system and network monitoring application, watching hosts and services specified by the administrator, and providing alerts when problems emerge or when problems subside. Some key features of *Nagios* include:

- Monitoring of network services such as SMTP
- Monitoring of host resources such as processor load
- Simple plugin design that allows administrators to easily develop their own service checks
- Parallelized service checks
- Contact notifications when service or host problems occurs and get resolved
- Ability to define event handlers to be run during service or host events for proactive problem resolution
- Support for implementing redundant monitoring hosts

Optional web interface for viewing current network status, notification and problem history, log file, and so on.
Nagios is used to report information collected via the infrastructure of *SuperMon*. The data collected by *SuperMon* includes both system performance metrics like CPU and memory utilization, as well as environmental data such as fan speed and temperature. This data is collected on a regular basis via a call to the *SuperMon* daemon on the head node, which communicates with daemons on the other nodes. The *Nagios* plug-in requests the data, causes it to be entered in the CMDB, and then displays the result. The following is a partial list of node-health and processor-health metrics: context switches per second, interrupts per second, I/O in/out of kernel per second, 1-minute, 5-minute and 15-minute processor load averages, free or shared memory, memory used as buffers or cache, processor-usage (idle, nice, system, user), number of processes, condition of realtime clock battery, swap free space, fan speed (if available), temperature (if available), processor clock speed, In addition to the real-time collection of events, the XC Software also provides a scalable logging facility, *syslog-ng*. *Syslog-ng* can coexist or totally replace the traditional syslog functionality. It provides more flexibility to handle logs, adds better filters, and contains a better forwarding mechanism. In order to control the size of global system logs, events can be prioritized so that only the highest priority events are forwarded to a centralized log file maintained on the head nodes.

The XC System is typically managed from the head node using the XC Management Network. One of the key system-control utilities is *Parallel Distributed Shell* (pdsh). Pdsh is a multithreaded remote shell client used to execute commands on multiple remote hosts in parallel. The implementation of pdsh relies on the use of ssh as the main remote shell service. Separate facilities are provided on the head node, for access to console commands. Subsequent to initial configuration, the administrator can perform and log console commands such as starting, stopping, and booting nodes from the head node communicating over the console branch of the XC Administrative Network (the XC Console Network). For Cluster Platform 3000 (Intel Xeon) and Cluster Platform 6000 systems (Intel Itanium), this is implemented with the XC Console Management Facility (CMF) which was adopted from HP AlphaServer SC System Software, and depends on the fact that the XC Console Network is connected to the remote management port on each of the nodes. For Cluster Platform 4000 systems (AMD Opteron), the console connections are made to the management port on each node and the administrator performs console operations using telnet. Console-logging is not available for the Cluster Platform 4000 systems.

Standard Features

Serviceability

XC System Software includes some core components which facilitate the serviceability of the XC System and the XC System Software. For detection of potential problems relating to the XC System Interconnect, there are interconnect-specific diagnostic tools as well as generic tests which stress the network. An XC Installation and Operational Verification procedure can be used to check whether system installation has been done correctly, and whether the system behaves correctly. During normal operation of the XC System, various facilities are provided to help detect and diagnose problems. Crash dump support is provided as a way to trace the origins of software errors (which, in some cases, could be triggered by hardware failures). The sys_check utility is able to collect data and log files to further facilitate problem analysis. Hardware failures can be anticipated or detected by environmental monitoring using Nagios and Supermon as described elsewhere in this document.

Availability

For increased redundancy and availability, the XC System may be configured with the administration and login services available on two different nodes (administration nodes). The XC System may also include separate I/O nodes with shared access to SAN storage. Should a node fail running one of these services (administration, login or I/O), the service can be manually resumed on its alternate administration node. In addition, there are some key redundancy features at the application level that help to limit single points of failure. The *SLURM* and LSF job launch mechanisms, as well as scheduling and resource management, provide application level failover support between the two administration nodes. For example, the *SLURM* control daemon can be optionally configured with a backup to allow failover in the event of a node failure. If the administration node running the *SLURM* and LSF services fails, this failure will be detected. The services will then resume on the alternate administration node (e.g. an LSF daemon will be started on that node) with no disruption to users, or to pending jobs. Running jobs will be restarted.

Modules

The XC System Software includes the *Modules* package (not to be confused with kernel modules), which provides for the dynamic modification of a user's environment via module files. The module command processes the information contained in the module file to alter or set the shell environment variables, such as PATH and MANPATH. Users can add their own module files to further modify their environment. The module files that are automatically installed and that are specific to XC System Software include Intel-compilers, mpi, and modules.

MPI

XC System Software includes HP MPI V2.1.1 (HP MPI), and is fully integrated with *SLURM* resource management including job-launch as described above. HP MPI complies with the MPI-1.2 and MPI-2 standards and is a high-performance, robust, high-quality, native implementation. HP MPI uses OS-specific enhancements to provide low latency and high bandwidth point-to-point and collective communication routines. HP MPI supports 32- and 64-bit applications, single- and multi-threaded, and provides tools to debug and instrument MPI execution. Also offered is MPICH 1.2.5 compatibility when an application is built shared with HP MPI. Included in HP MPI is support for Quadrics Ltd. QsNet^{II}® and Gigabit Ethernet on Integrity systems and Myrinet GM2.1 and Gigabit Ethernet on HP ProLiant, including standard TCP/IP support. Applications built with HP MPI are interconnect-independent; it is not necessary to recompile or relink an application in order for it to operate on a different XC System Interconnect.

Standard Features

MLIB

XC System Software includes HP MLIB V1.5 (HP MLIB) for compatibility with previous versions of the XC System Software, but HP MLIB will be removed from subsequent releases of the XC System Software. HP MLIB is one of several mathematical libraries for Linux systems, which support application developers looking for ways to speed up development of new applications and to shorten the execution time of long-running technical applications. HP MLIB contains robust callable subprograms, including all BLAS 1, 2, and 3 subroutines, sparse BLAS subroutines, and a collection of commonly used dense and sparse linear system solvers, including LAPACK, ScaLAPACK, SuperLU_DIST, Fast Fourier Transforms (FFTs), and convolutions. Many of these capabilities are included in math libraries from Intel (MKL), AMD (ACML), and other software vendors. These libraries are widely used today and therefore to encourage increasing levels of portability and standardization, the use of HP MLIB is being deprecated in favor of the more popular mathematical libraries.

Linux

XC System Software includes XC Linux for High Performance Computing (HPC Linux), an LSB-compliant Gnu/Linux distribution (kernel and supporting software packages) based on Linux Kernel Version 2.4. Linux is the open source operating system kernel, created by Linus Torvalds in the early 1990s, and subsequently adopted and maintained by thousands of people world-wide. All official releases of the Linux kernel can be found at the kernel.org web site, http://www.kernel.org. Hewlett-Packard will periodically update the XC System Software with new versions of the Linux kernel or with patches to the Linux kernel (see the section on Software Product Services). Any changes made to the kernel, other than those provided or authorized by Hewlett-Packard, will invalidate all HP software support contracts for XC System Software.

The system utilities and services come from the Gnu project and from a variety of other open source projects. The base collection of Gnu/Linux utilities provided with the XC System Software is derived from the Red Hat Enterprise Linux Advanced Server 3.0 packages of Red Hat, Inc. The set of packages included with the XC System Software is neither supported nor endorsed by Red Hat. However, these packages are covered by HP software support contracts for XC System Software, provided that no changes are made to those packages, except as provided or explicitly authorized by Hewlett-Packard. Any applications that run correctly using Red Hat Enterprise Linux Advanced Server Version 3.0 will also run correctly using HPC Linux. The following supported packages (RPMs) are included in the XC System Software.

Linux Supported Packages

4Suite-0.11.1-14
a2ps-4.13b-28
acl-2.2.3-1
acpid-1.0.2-2
alchemist-1.0.27-1
ant-1.5.2-23
ant-libs-1.5.2-23
arts-1.1.3-2.3
ash-0.3.8-16
aspell-0.33.7.1-25.3
aspell-config-0.33.7.1-25.3
aspell-devel-0.33.7.1-25.3
at-3.1.8-60_EL3
atk-1.2.4-3.0
at-spi-1.1.9-1
attr-2.2.0-1
audiofile-0.2.3-7.1
authconfig-4.3.7-1
authconfig-gtk-4.3.7-1
authd-1.4.1-1.rhel3
autoconf-2.57-3
autoconf213-2.13-6
autofs-4.1.3-47
automake-1.6.3-5
hptc-pdsh-support-1-0
hptc-power-2.2-2
hptc-qsnet2-diag-1-5
hptc-slurm-0.21-1hp
hptc-supermon-1-4.8
hptc-supermon-config-1-5
hptc-supermon-modules-1-4.k2.4.21_27.5hp.XCsmp
hptc-syslogng-1-3
hptc-sysman-1-1.11
hptc-sysmandb-1-16
htmlview-2.0.0-10
httpd-2.0.46-44.LNXHPC
hwbrowser-0.11-1
hwdata-0.101.14-1
ibhost-biz-3.0.16.5_2-2hptc.k2.4.21_27.5hp.XCsmp
imlib-1.9.13-13.3
im-sdk-20030118-6
indent-2.2.9-2
indexhtml-3-2.LNXHPC
info-4.5-3
initscripts-7.31.18.EL-1.LNXHPC
intltool-0.26-1
pam-0.75-62
pam-0.75-62
pam-devel-0.75-62
pango-1.2.5-5
parted-1.6.3-29.3
passwd-0.68-3.1
patch-2.5.4-16
patchutils-0.2.23-1
pax-3.0-6
pciutils-2.1.10-7
pciutils-devel-2.1.10-7
pcre-3.9-10.1
pdksh-5.2.14-21
pdsh-1.7-6
perl-5.8.0-88.9
perl-AppConfig-1.52-4
perl-CGI-2.81-88.9
perl-CPAN-1.61-88.9
perl-DateManip-5.42a-0.rhel3
perl-DBD-MySQL-2.1021-3
perl-DBI-1.32-9
perl-Filter-1.29-3
perl-HTML-Parser-3.26-17
perl-HTML-Tagset-3.03-28

Standard Features

automake14-1.4p6-6
automake15-1.5-7
basesystem-8.0-2
bash-2.05b-29.0.3
bc-1.06-15
bcel-5.0-10
beecrypt-3.0.1-0.20030630
bind-9.2.4-5_EL3
bind-libs-9.2.4-5_EL3
bind-utils-9.2.4-5_EL3
binutils-2.14.90.0.4-35
bison-1.875c-2.EL
bitmap-fonts-0.3-2
bonobo-activation-2.2.2-1
byacc-1.9-25
bzip2-1.0.2-11
bzip2-devel-1.0.2-11
bzip2-libs-1.0.2-11
cdecl-2.5-27
cdparanoia-libs-alpha9.8-15
cdrecord-2.01.0.a32-0.EL3.1
chkconfig-1.3.11-0.3
chkfontpath-1.9.10-1.RHEL
commons-beanutils-1.6.1-10
commons-collections-2.1-9
commons-digester-1.4.1-10
commons-logging-1.0.2-12
commons-modeler-1.0-5
compat-db-4.0.14-5.1
compat-db-4.0.14-5.1
compat-gcc-7.3-2.96.128
compat-gcc-c++-7.3-2.96.128
compat-libstdc++-7.3-2.96.128
compat-libstdc++-devel-7.3-2.96.128
comps-extras-9.0.4-1
control-center-2.2.0.1-13
coreutils-4.5.3-26
cpio-2.5-3
cpp-3.2.3-49
cproto-4.6-16
cracklib-2.7-22
cracklib-2.7-22
cracklib-dicts-2.7-22
crash-3.8-3
crontabs-1.10-5
crypto-utils-2.0-4
ctags-5.4-2
cups-1.1.17-13.3.24
cups-libs-1.1.17-13.3.24
cups-libs-1.1.17-13.3.24
cup-v10k-10
curl-7.10.6-4.2
curl-devel-7.10.6-4.2
cvs-1.11.2-25
IO-Socket-SSL-0-96
iproute-2.4.7-11.30E.1
ipsec-tools-0.2.5-0.6
iptables-1.2.8-12.3
iptables-ipv6-1.2.8-12.3
iputils-20020927-11.30.1
ipvsadm-1.21-9.ipvs108
irda-utils-0.9.15-1
isdn4k-utils-3.1-76
jaf-20030319-1
jakarta-regexp-1.2-12
javamail-20031006-1
jfsutils-1.1.2-2
jpackage-utils-1.5.38-1jpp_4rh
junit-3.8.1-1
jwhois-3.2.2-1
kbd-1.08-10.2
kdebase-3.1.3-5.5
kdelibs-3.1.3-6.6.LNXHPC
kernel-2.4.21-27.0.2.EL.1.LNXHPC
kernel-smp-2.4.21-27.5hp.XC
kernel-utils-2.4-8.37.7
kickstart-1-9
krb5-devel-1.2.7-38
krb5-libs-1.2.7-38
krb5-libs-1.2.7-38
krb5-workstation-1.2.7-38
krbafs-1.1.1-11
krbafs-1.1.1-11
krbafs-devel-1.1.1-11
krbafs-utils-1.1.1-11
kudzu-1.1.22.9-1
kudzu-devel-1.1.22.9-1
laus-0.1-66RHEL3
laus-libs-0.1-66RHEL3
laus-libs-0.1-66RHEL3
less-378-12
lftp-2.6.3-5
lha-1.14i-10.4
libacl-2.2.3-1
libacl-2.2.3-1
libaio-0.3.96-5
libaio-devel-0.3.96-5
libao-0.8.3-3.1
libao-devel-0.8.3-3.1
libart_lgpl-2.3.11-2
libattr-2.2.0-1
libattr-2.2.0-1
libbonobo-2.2.3-1
libbonoboui-2.2.2-1
libcap-1.10-15.1
libcap-1.10-15.1
libcap-devel-1.10-15.1
libdbi-0.6.5-5
perl-libwww-perl-5.65-6
perl-libxml-enno-1.02-29
perl-libxml-perl-0.07-28
perl-Parse-Yapp-1.05-30
perl-URI-1.21-7
perl-XML-Dumper-0.4-25
perl-XML-Encoding-1.01-23
perl-XML-Grove-0.46alpha-25
perl-XML-Parser-2.31-15
perl-XML-Simple-1.08-1
perl-XML-Twig-3.09-3
php-4.3.2-19.ent
php-imap-4.3.2-19.ent
php-ldap-4.3.2-19.ent
pilot-link-0.11.7-1.1
pinfo-0.6.6-4
pkgconfig-0.14.0-5
pnm2ppa-1.04-7
popt-1.8.2-13
popt-1.8.2-13
portmap-4.0-56
ppp-2.4.1-14.1
prelink-0.3.2-2.EL
printman-0.0.1-0.20021202.13
procmail-3.22-9
procps-2.0.17-13
psacct-6.3.2-28.rhel3
psmisc-21.3-1.RHEL.0
pspell-0.12.2-16.1
pspell-devel-0.12.2-16.1
psutils-1.17-21
pwlib-1.4.7-7.EL
pygtk2-1.99.16-8
pygtk2-libglade-1.99.16-8
pyorbit-1.99.3-5
python-2.2.3-6
python-devel-2.2.3-6
pyxf86config-0.3.5-1
PyXML-0.7.1-9
qsnet2diags-1.0.12-3.1hptc
qsnet2diagscommon-1.0.12-2.1hptc
qsnet2libs-1.8.13-0.1hptc
qsnetdiags-1.0.2-14.1hptc
qsnetlibs-1.4.30-1.1hptc
qsqsql-1.0.12-2.2hptc
qsswm-2.0.7-0.1hptc
qt-3.1.2-13.4
quota-3.10-4
raidtools-1.00.3-7
rcs-5.7-20
rdate-1.3-2
rdist-6.1.5-35.30.1
readline-4.3-5.2
readline-4.3-5.2

Standard Features

cyrus-sasl-2.1.15-10
cyrus-sasl-2.1.15-10
cyrus-sasl-devel-2.1.15-10
cyrus-sasl-gssapi-2.1.15-10
cyrus-sasl-md5-2.1.15-10
cyrus-sasl-plain-2.1.15-10
db4-4.1.25-8
db4-4.1.25-8
db4-devel-4.1.25-8
db4-utils-4.1.25-8
ddd-3.3.1-23
dejagnu-1.4.2-10
desktop-backgrounds-basic-2.0-14
desktop-backgrounds-extra-2.0-14
desktop-file-utils-0.3-10
desktop-printing-0.1.10-15.1
dev-3.3.12.3-1
devlabel-0.48.03-6
dhclient-3.0.1-10_EL3
dhcp-3.0.1-10_EL3
dialog-0.9b-20020814.6
diffstat-1.31-2
diffutils-2.8.1-8
diskdumputils-0.4.0-1
distcache-0.4.2-9.3
distcache-devel-0.4.2-9.3
docbook-dtds-1.0-17.2
dos2unix-3.1-15
dosfstools-2.8-10
doxygen-1.3.5-0.RHEL3.1
dump-0.4b28-7
e2fsprogs-1.32-15.1
e2fsprogs-devel-1.32-15.1
eal3-certification-1.1-2
eal3-certification-doc-1.1-2
ed-0.2-33
eel2-2.2.4-3.E
eject-2.0.13-2
ElectricFence-2.2.2-15
elfutils-0.91-3
elfutils-devel-0.91-3
elfutils-libelf-0.91-3
elfutils-libelf-0.91-3
elfutils-libelf-devel-0.91-3
elinks-0.4.2-7
emacs-21.3-4
enscript-1.6.1-24.4
eog-2.2.2-1
esound-0.2.28-6
ethtool-1.8-3.3
evolution-1.4.5-9
evolution-connector-1.4.7-5
expat-1.95.5-6
expat-1.95.5-6
libdbi-dbd-mysql-0.6.5-5
libelanhosts-0.9-1.3hp
libf2c-3.2.3-49
libgail-gnome-1.0.2-1
libgal2-1.99.10-1
libgcc-3.2.3-49
libgcc-3.2.3-49
libgcc-ssa-3.5ssa-0.20030801.48
libgcj-3.2.3-49
libgcj-3.2.3-49
libgcj-devel-3.2.3-49
libgcj-ssa-3.5ssa-0.20030801.48
libgcj-ssa-devel-3.5ssa-0.20030801.48
libghttp-1.0.9-8.1
libglade2-2.0.1-3
libgnat-3.2.3-49
libgnome-2.2.2-6
libgnomecanvas-2.2.0.2-2
libgnomeprint22-2.2.1.3-5
libgnomeprintui22-2.2.1.3-2.0
libgnomeui-2.2.1-1
libgsf-1.6.0-4
libgtop2-2.0.2-3
libIDL-0.8.0-9
libjpeg-6b-30
libjpeg-6b-30
libmng-1.0.4-3
libmudflap-3.5ssa-0.20030801.48
libmudflap-devel-3.5ssa-0.20030801.48
libobjc-3.2.3-49
libobjc-3.2.3-49
libogg-1.0-5.1
libogg-devel-1.0-5.1
libol-0.3.13-1
libol-devel-0.3.13-1
libpng-1.2.2-25
libpng-1.2.2-25
libpng10-1.0.13-15
libraw1394-0.9.0-10.RHEL3
librsvg2-2.2.3-6
libsoup-1.99.26-1
libstdc++-3.2.3-49
libstdc++-3.2.3-49
libstdc++-devel-3.2.3-49
libstdc++-ssa-3.5ssa-0.20030801.48
libstdc++-ssa-devel-3.5ssa-0.20030801.48
libtermcap-2.0.8-35
libtermcap-2.0.8-35
libtermcap-devel-2.0.8-35
libtiff-3.5.7-22.el3
libtiff-3.5.7-22.el3
libtool-1.4.3-6
libtool-libs-1.4.3-6
readline-devel-4.3-5.2
redhat-artwork-0.73.2-1.LNXHPC
redhat-config-bind-2.0.0-14.2
redhat-config-date-1.5.22-3
redhat-config-httpd-1.1.0-4
redhat-config-keyboard-1.1.4-1
redhat-config-kickstart-2.3.22-3
redhat-config-language-1.0.14-2
redhat-config-mouse-1.0.13-1
redhat-config-netboot-0.1.5-1
redhat-config-network-1.2.63-1
redhat-config-network-tui-1.2.63-1
redhat-config-nfs-1.0.13-6
redhat-config-packages-1.2.9.2-1
redhat-config-printer-0.6.47.3.23-1
redhat-config-printer-gui-0.6.47.3.23-1
redhat-config-proc-0.23-0.EL3.1
redhat-config-rootpassword-1.0.6-1
redhat-config-samba-1.0.16-2
redhat-config-securitylevel-1.2.9.2-1
redhat-config-securitylevel-tui-1.2.9.2-1
redhat-config-services-0.8.5-19.RHEL3.1
redhat-config-soundcard-1.0.8-1
redhat-config-users-1.1.18-2
redhat-config-xfree86-0.7.7-1.LNXHPC
redhat-java-rpm-scripts-1.0.2-2
redhat-logviewer-0.9.3-5
redhat-lsb-1.3-3.1.EL3
redhat-lsb-1.3-3.1.EL3
redhat-menus-0.39-1
redhat-rpm-config-8.0.28-2
redhat-switch-mail-0.5.20-1
redhat-switch-mail-gnome-0.5.20-1
rhgb-0.11.2-4
rhpl-0.110.6-1
rh-postgresql-devel-7.3.8-2
rh-postgresql-libs-7.3.8-2
rmt-0.4b28-7
rootfiles-7.2-6
rpm-4.2.3-13
rpm-build-4.2.3-13
rpm-devel-4.2.3-13
rpm-libs-4.2.3-13
rpm-python-4.2.3-13
rp-pppoe-3.5-4.1
rsh-0.17-17
rsync-2.6.2-1
rsync-2.6.3-1.LNXHPC
rusers-0.17-40.30.2
rwho-0.17-18
samba-3.0.9-1.3E.2
samba-3.0.9-1.3E.2
samba-client-3.0.9-1.3E.2
samba-common-3.0.9-1.3E.2

Standard Features

expat-devel-1.95.5-6
expect-5.38.0-92.2
fam-2.6.8-15
fam-devel-2.6.8-15
fbset-2.1-13
fetchmail-6.2.0-3
file-3.39-9
file-roller-2.2.3-6
filesystem-2.2.1-3
findutils-4.1.7-9
finger-0.17-18
finger-server-0.17-18
firstboot-1.1.31-1.LNXHPC
flamethrower-0.1.6-1
flex-2.5.4a-29
fontconfig-2.2.1-13
fontconfig-2.2.1-13
fontilus-0.3-4
foomatic-2.0.2-15.1.1
freetype-2.1.4-4.0
freetype-2.1.4-4.0
ftp-0.17-17
gail-1.2.2-1
gaim-1.0.1-1.RHEL3
gawk-3.1.1-9
gcc-3.2.3-49
gcc-c++-3.2.3-49
gcc-c++-ssa-3.5ssa-0.20030801.48
gcc-g77-3.2.3-49
gcc-g77-ssa-3.5ssa-0.20030801.48
gcc-gnat-3.2.3-49
gcc-java-3.2.3-49
gcc-java-ssa-3.5ssa-0.20030801.48
gcc-objc-3.2.3-49
gcc-objc-ssa-3.5ssa-0.20030801.48
gcc-ssa-3.5ssa-0.20030801.48
GConf2-2.2.1-4
gconf-editor-0.5.0-1
gd-1.8.4-12.3.1
gdb-6.1post-1.20040607.52
gdbm-1.8.0-20
gdbm-1.8.0-20
gdbm-devel-1.8.0-20
gdk-pixbuf-0.22.0-11.3.3
gdm-2.4.1.6-5
gedit-2.2.2-1
gettext-0.11.4-7
gftp-2.0.14-2
ggv-2.0.1-4
ghostscript-7.05-32.1.9
ghostscript-fonts-5.50-9
gimp-print-4.2.4-5.1
gimp-print-utils-4.2.4-5.1
glib-1.2.10-11.1
libungif-4.1.0-15
libusb-0.1.6-3
libusb-devel-0.1.6-3
libuser-0.51.7-1.EL3.3
libuser-devel-0.51.7-1.EL3.3
libvorbis-1.0-7
libvorbis-devel-1.0-7
libwnck-2.2.1-3
libwvstreams-3.70-10
libxml-1.8.17-9.1
libxml2-2.5.10-7
libxml2-devel-2.5.10-7
libxml2-python-2.5.10-7
libxslt-1.0.33-5
linc-1.0.3-1
lockdev-1.0.1-1.2
lockdev-devel-1.0.1-1.2
logrotate-3.6.9-1
logwatch-4.3.2-2
losetup-2.11y-31.2
lsf-6.0-13.hptc
lslk-1.29-8
lsof-4.63-4
ltrace-0.3.32-3.EL
lvm-1.0.8-9
m3-dist-1.0.14-1.1hptc
m4-1.4.1-13
magicdev-1.1.4-8.0
mailcap-2.1.14-1
mailx-8.1.1-31
make-3.79.1-17
MAKEDEV-3.3.12.3-1
man-1.5k-10
man-pages-1.60-4.2
mdadm-1.5.0-9
metacity-2.4.55-7.12
mgetty-1.1.30-3
mikmod-3.1.6-21.2
mingetty-1.06-1
minicom-2.00.0-17.1
mkinitrd-3.5.13-2.LNXHPC
mkisofs-2.01.0.a32-0.EL3.1
mktemp-1.5-18
MLIB-XC3000-1.5-2
MLIB-XC4000-1.5-4
mod_authz_ldap-0.22-5
mod_perl-1.99_09-10.ent
mod_python-3.0.3-3.ent
mod_ssl-2.0.46-44.LNXHPC
modulefiles_hptc-1.0.4-041006
modules-3.1.6-2hptc
modutils-2.4.25-14.EL
modutils-devel-2.4.25-14.EL
mount-2.11y-31.2
samba-swat-3.0.9-1.3E.2
schedutils-1.3.0-5
scrollkeeper-0.3.11-8E
SDL-1.2.5-6
sed-4.0.7-3
sendmail-8.12.11-4.RHEL3.1
setarch-1.3-1
setserial-2.17-12
setup-2.5.27-1.LNXHPC
setuptool-1.13-1
sgml-common-0.6.3-14
shadow-utils-4.0.3-22.02
sharutils-4.2.1-16
slang-1.4.5-18
slang-devel-1.4.5-18
slocate-2.7-3
slurm-0.3.8-1.36hp
slurm-auth-munge-0.3.8-1.36hp
slurm-auth-none-0.3.8-1.36hp
slurm-devel-0.3.8-1.36hp
slurm-sched-wiki-0.3.8-1.36hp
slurm-switch-elan-0.3.8-1.36hp
sox-12.17.4-4.3
specspo-3EL-1
splint-3.0.1.7-0.20030123
squid-2.5.STABLE3-6.3E.2
star-1.5a08-4
startup-notification-0.5-1
strace-4.5.6-1.EL
stunnel-4.04-4
sudo-1.6.7p5-1
swig-1.1p5-22
switchdesk-3.9.8-17
switchdesk-gnome-3.9.8-17
symlinks-1.2-18
sys_check-1.0.2-7
sysklogd-1.4.1-12.3
syslinux-2.06-0.3E
syslog-ng-1.6.2-1
sysreport-1.3.7.2-2
systemconfigurator-2.0.9-hp6
systemimager-client-3.2.3-8hp
systemimager-common-3.2.3-8hp
systemimager-flamethrower-3.2.3-8hp
systemimager-server-3.2.3-8hp
systemimager-x86_64boot-standard-3.2.3-8hp
SysVinit-2.85-4.2
talk-0.17-20
tar-1.13.25-13
tcl-8.3.5-92.2
tcl-8.3.5-92.2
tclx-8.3-92.2
tcp_wrappers-7.6-34.1

Standard Features

glib-1.2.10-11.1
glib2-2.2.3-2.0
glibc-2.3.2-95.30
glibc-2.3.2-95.30
glibc-common-2.3.2-95.30
glibc-devel-2.3.2-95.30
glibc-devel-2.3.2-95.30
glibc-headers-2.3.2-95.30
glibc-kernheaders-2.4-8.34.1
glibc-profile-2.3.2-95.30
glibc-utils-2.3.2-95.30
gm-2.1.7_Linux-2.1hptc
gmp-4.1.2-5
gmp-devel-4.1.2-5
gnome-applets-2.2.2-2
gnome-audio-1.4.0-6
gnome-desktop-2.2.2-1.LNXHPC
gnome-icon-theme-1.0.5-1.E
gnome-libs-1.4.1.2.90-34.2
gnome-media-2.2.1.1-4
gnomemeeting-0.96.0-4
gnome-mime-data-2.2.0-3
gnome-panel-2.2.2.1-8.LNXHPC
gnome-pilot-2.0.10-1
gnome-python2-1.99.14-5
gnome-python2-applet-1.99.14-5
gnome-python2-bonobo-1.99.14-5
gnome-python2-canvas-1.99.14-5
gnome-python2-gtkhtml2-1.99.14-5
gnome-session-2.2.2-3.LNXHPC
gnome-spell-1.0.5-1
gnome-system-monitor-2.0.5-2
gnome-terminal-2.2.2-1
gnome-themes-2.2-3.1
gnome-user-docs-2.0.1-3
gnome-utils-2.2.3-1
gnome-vfs2-2.2.5-2E.1
gnome-vfs2-extras-0.99.10-3
gperf-2.7.2-9
gpm-1.19.3-27.2
gpm-1.19.3-27.2
gpm-devel-1.19.3-27.2
grep-2.5.1-24.1
groff-1.18.1-27
grub-0.93-4.3.1.LNXHPC
gstreamer-0.6.0-5
gstreamer-plugins-0.6.0-14
gstreamer-tools-0.6.0-5
gtk+-1.2.10-31
gtk2-2.2.4-12
gtk2-engines-2.2.0-2
gtk-engines-0.12-2.0
gtkhtml2-2.2.0-6
gtkhtml3-3.0.9-7
mozilla-1.4.3-3.0.7.1.LNXHPC
mozilla-chat-1.4.3-3.0.7.1.LNXHPC
mozilla-mail-1.4.3-3.0.7.1.LNXHPC
mozilla-nspr-1.4.3-3.0.7.1.LNXHPC
mozilla-nss-1.4.3-3.0.7.1.LNXHPC
mpage-2.5.3-3
mtools-3.9.8-8
mtr-0.52-2
mt-st-0.7-11
munge-0.3-1.6hp
mute-1.9.6-1.1hptc
mutt-1.4.1-3.4
mx-2.0.3-8
mx4j-1.1.1-6
MyODBC-2.50.39-12
mysql-3.23.58-2.3
mysql-devel-3.23.58-2.3
MySQL-Max-4.0.20-0
MySQL-python-0.9.1-6
MySQL-server-4.0.20-0
nagios-bundle-1-1.2hptc
nano-1.2.1-4
nautilus-2.2.4-4.E.LNXHPC
nautilus-cd-burner-0.3.2-2E
nautilus-media-0.2.1-5.1E
nc-1.10-18
ncompress-4.2.4-38
ncurses-5.3-9.3
ncurses-5.3-9.3
ncurses-devel-5.3-9.3
Net_SSLeay.pm-1-25
netconfig-0.8.19-1.1
netdump-0.6.11-3
netdump-server-0.6.11-3
net-snmp-5.0.9-2.30E.12
net-snmp-devel-5.0.9-2.30E.12
net-snmp-libs-5.0.9-2.30E.12
net-snmp-utils-5.0.9-2.30E.12
net-tools-1.60-20E.1
newt-0.51.5-1
newt-devel-0.51.5-1
newt-perl-1.08-4
nfs-utils-1.0.6-33EL
nptl-devel-2.3.2-95.30
nscd-2.3.2-95.30
nss_db-2.2-20.4
nss_db-2.2-20.4
nss_ldap-207-11
nss_ldap-207-11
ntp-4.1.2-4.EL3.1
ntsysv-1.3.11-0.3
Omni-0.7.2-4
Omni-foomatic-0.7.2-4
openh323-1.11.2-4.1E
tcpdump-3.7.2-7.E3.2
tcsh-6.12-11.EL3
telnet-0.17-26
termcap-11.0.1-17.1
texinfo-4.5-3
tftp-0.39-0.EL3.1
tftp-server-0.39-0.EL3.1
time-1.7-23
tk-8.3.5-92.2
tmpwatch-2.8.4-5
traceroute-1.4a12-20
ttfprint-0.9-8
ttmkfdir-3.0.9-6
tux-3.2.19-2
tzdata-2004e-1.EL
udpcast-20040531-1
unix2dos-2.2-19
unixODBC-2.2.8-2.3.0.2
unzip-5.50-34
urw-fonts-2.1-5.0
usbutils-0.11-1
usermode-1.68-5E.2
usermode-gtk-1.68-5E.2
utempter-0.5.5-1.3EL.0
util-linux-2.11y-31.2
vconfig-1.6-2
VFlib2-2.25.6-17
vim-common-6.3.046-0.30E.1
vim-minimal-6.3.046-0.30E.1
vixie-cron-3.0.1-75.1
vnc-server-4.0-0.beta4.1.4
vsftpd-1.2.1-3E.1
vte-0.11.10-11
webalizer-2.01_10-15.ent
wget-1.8.2-15
which-2.14-7
wireless-tools-26-2
words-2-21
wvdial-1.53-11
xalan-j-2.4.1-11
Xaw3d-1.5-18
xchat-2.0.4-4.EL
xerces-j-2.2.1-11
XFree86-100dpi-fonts-4.3.0-78.EL
XFree86-4.3.0-78.EL
XFree86-75dpi-fonts-4.3.0-78.EL
XFree86-base-fonts-4.3.0-78.EL
XFree86-doc-4.3.0-78.EL
XFree86-font-utils-4.3.0-78.EL
XFree86-libs-4.3.0-78.EL
XFree86-libs-4.3.0-78.EL
XFree86-libs-data-4.3.0-78.EL
XFree86-Mesa-libGL-4.3.0-78.EL
XFree86-Mesa-libGL-4.3.0-78.EL

Standard Features

gzip-1.3.3-9
hdparm-5.4-1
hesiod-3.0.2-28.1
hesiod-devel-3.0.2-28.1
hotplug-2002_04_01-20.2
hpasm-7.2.0b-204
hp-collectl-1.3.1-1
hpc-release-3.0-6.LNXHPC
hpijs-1.3-32.1.9
hpls-pwr-2.1-0hptc
hpmpi-2.01.01-02
hpoj-0.90-14
hptc_ovp-1.6-12
hptc_release-1.0-3
hptc-3rd-party-perl-mods.pm-1-1
hptc-cluconfig-2-12
hptc-cmf-1-4.0
hptc-dgemm-1-1
hptc-gm-diag-1-4
hptc-hpasm-config-1-2
hptc-hpasm-post-1-0
hptc-ire-1.1.0-14
hptc-lm-1.0-6
hptc-lsf-0.1-17hp
hptc-md5deep-1-2
hptc-nagios-config-1-1.3
hptc-nconfig-1.0-22
openjade-1.3.1-12.1
openldap-2.0.27-17
openldap-2.0.27-17
openldap-clients-2.0.27-17
openldap-devel-2.0.27-17
openldap-servers-2.0.27-17
openmotif-2.2.3-3.RHEL3
openmotif-2.2.3-3.RHEL3
openssh-3.6.1p2-33.30.3
openssh-askpass-3.6.1p2-33.30.3
openssh-askpass-gnome-3.6.1p2-33.30.3
openssh-clients-3.6.1p2-33.30.3
openssh-server-3.6.1p2-33.30.3
openssl-0.9.7a-33.12
openssl-0.9.7a-33.12
openssl096b-0.9.6b-16
openssl-devel-0.9.7a-33.12
oprofile-0.5.4-22
oprofile-devel-0.5.4-22
ORBit-0.5.17-10.4
ORBit2-2.6.2-1
pam_krb5-1.73-1
pam_krb5-1.73-1
pam_passwdqc-0.7.5-1
pam_slurm-1.2.1-1.2hptc
pam_smb-1.1.7-1
pam_smb-1.1.7-1
XFree86-Mesa-libGLU-4.3.0-78.EL
XFree86-Mesa-libGLU-4.3.0-78.EL
XFree86-tools-4.3.0-78.EL
XFree86-truetype-fonts-4.3.0-78.EL
XFree86-twm-4.3.0-78.EL
XFree86-xauth-4.3.0-78.EL
XFree86-xdm-4.3.0-78.EL
XFree86-xfs-4.3.0-78.EL
XFree86-Xnest-4.3.0-78.EL
XFree86-Xvfb-4.3.0-78.EL
xinetd-2.3.12-6.3E
xinitrc-3.32-1.LNXHPC
xloadimage-4.1-27
xml-common-0.6.3-14
xpdf-2.02-9.5
xscreensaver-4.10-8.LNXHPC
xsri-2.1.0-5.LNXHPC
xterm-179-5.EL
yelp-2.2.3-1.E
ypbind-1.12-5.21.1
ypserv-2.8-6.1hp
yp-tools-2.8-6
zip-2.3-16
zlib-1.1.4-8.1
zlib-1.1.4-8.1
zlib-devel-1.1.4-8.1

Hardware Requirements

Overview of Processors and Systems Supported

XC System Software must be run on each node of a valid configuration, such as an HP Cluster Platform 3000, an HP Cluster Platform 4000 or an HP Cluster Platform 6000. The supported Cluster Platform 3000 systems are based on HP ProLiant DL140 G2, DL360 G4 and DL 380 G4 servers each of which use Intel® Xeon™ processors with Intel® Extended Memory 64 Technology. The supported Cluster Platform 4000 systems are based on HP ProLiant DL145 G1, DL145 G2, DL385 G1 and DL585 G1 servers each of which use AMD Opteron™ processors. The supported Cluster Platform 6000 systems are based on HP Integrity rx1620, rx2600 and rx2620 systems with Intel® Itanium® processors. Each node must be connected to the others with a valid instance of an XC System Interconnect, a Cluster Platform Administrative Network, and a Cluster Platform Console Network. XC Systems with 32 nodes or less, and which have an XC System Interconnect consisting of Gigabit Ethernet, may be configured to combine the XC System Interconnect and the Cluster Platform Administrative Network into a single network with a single Gigabit Ethernet switch. The supported XC System Interconnects for both Cluster Platform 3000 and Cluster Platform 4000 systems are Gigabit Ethernet, Infiniband and Myricom Myrinet-2000. In addition, the Cluster Platform 4000 also supports, as the XC System Interconnect, Quadrics QsNet^{II}. The supported XC System Interconnects for Cluster Platform 6000 systems are Gigabit Ethernet, Infiniband, and QsNet II. (XC System Software is also supported on certain custom-designed Integrity clusters with QsNet^{II} QsNet [Elan3].) Valid hardware configurations and options must comply with the XC System specification provided in the *HP Cluster Platform 3000, HP Cluster Platform 4000 or HP Cluster Platform 6000 Documentation CD-ROMs (part numbers AD162A, AD163A, and AD164A)* or that are electrically equivalent to such a specification. Any other hardware configuration or option may be considered invalid unless the software documentation explicitly states otherwise. XC System Software installation instructions can be found in the XC Installation Guide, which is part of the *XC System Software Administration Documentation Kit*.

Not all Cluster Platform configurations are supported by the XC System Software. Conversely, there may be customizable hardware configurations that are capable of being supported by the XC System Software. For more information on customizable configurations, please contact your HP sales representative. The following section describes the Cluster Platform configurations that are supported by XC System Software. For further details, consult the pertinent Cluster Platform documents in the *HP Cluster Platform 3000, HP Cluster Platform 4000 or HP Cluster Platform 6000 Documentation CD-ROMs (part numbers AD162A, AD163A, and AD164A)*.

Cluster Platform Administrative Network

The Cluster Platform Administrative Network is based on HP ProCurve 2848 and 2824 Gigabit Ethernet switches. Each Cluster Platform system has a single Utility Building Block (UBB) which includes a Cluster Platform Administrative Network root switch. In addition, each Cluster Platform may have several Compute Building Blocks (CBBs), each of which includes a Cluster Platform Administrative Network switch, connected to the Cluster Platform Administrative Network root switch in the UBB.

Cluster Platform Console Network

The Cluster Platform Console Network is based on HP ProCurve 2650 Ethernet 10/100 network switches. Each Cluster Platform system has a single Utility Building Block (UBB) which includes a Cluster Platform Console Network root switch. In addition, each Cluster Platform may have several Compute Building Blocks (CBBs), each of which includes a Cluster Platform Console Network switch, connected to the Cluster Platform Console Network root switch in the UBB. The Cluster Platform Console Network root switch is connected to the Cluster Platform Administrative Network root switch.

Configuration Information

XC System Interconnect

The XC System Interconnect provides high-speed connectivity for parallel applications. The XC System supports several different switch fabrics for use as an XC System Interconnect. The supported XC System Interconnects for both Cluster Platform 3000 and Cluster Platform 4000 systems are Gigabit Ethernet, Infiniband and Myricom Myrinet-2000. In addition, the Cluster Platform 4000 also supports, as the XC System Interconnect, Quadrics QsNet^{II}. The supported XC System Interconnects for Cluster Platform 6000 systems are Gigabit Ethernet, Infiniband and Quadrics QsNet (Elan 3) and QsNet^{II}(Elan 4). XC Systems with 32 nodes or less, and which have an XC System Interconnect consisting of Gigabit Ethernet, may be configured to combine the XC System Interconnect and the Cluster Platform Administrative Network into a single network with a single Gigabit Ethernet switch. All nodes of the XC System are directly attached to XC System Interconnect using one adapter per node. For the Quadrics-based systems, the adaptor is a Quadrics QM500 PCI-X Network Adapter. For the Myrinet-based systems, each XC System supports one of two possible adapters; the Myricom Myrinet-2000 PCI module Rev D or Rev E. For the Gigabit Ethernet-based systems, network connections either use onboard 1000Base-T NICs that come with certain of the servers, or separate adapters that are available as options to the servers. For Infiniband the adapter is a dual-port 4x HCA PCI-X adapter.

Nodes (Cluster Platform 3000)

The Cluster Platform 3000 consists of Proliant DL140 G2 and/or Proliant DL360 G4 and/or ProLiant DL380 G4 servers. One of the servers - often a ProLiant DL380 server - is designated as the head node. This node must have an internal disk drive, two processors, must have a DVD drive and must be attached to the Cluster Platform 3000's rack-mounted keyboard/monitor. The choice of a different head node is generally related to whether more PCI slots are required for extra connectivity to external systems.

Nodes (Cluster Platform 4000)

The Cluster Platform 4000 consists of ProLiant DL145 G1 and/or ProLiant DL145 G2 and/or ProLiant DL385 G1 and/or ProLiant DL585 G1 servers. One of the servers - usually a ProLiant DL385 server - is designated as the head node. The head node cannot be a ProLiant DL145 G1 server. The head node must have an internal disk drive, two processors, must have a DVD drive and must be attached to the Cluster Platform 4000's rack-mounted keyboard/monitor. The choice of a different head node is generally related to whether more PCI slots are required for extra connectivity to external systems.

Nodes (Cluster Platform 6000)

The Cluster Platform 6000 consists of Integrity rx1620 and/or Integrity rx2600 servers and/or Integrity rx2620 servers. One of the servers - often an rx2620 - is designated as the head node. This node must have an internal disk drive, two processors, must have a DVD drive and must be attached to the Cluster Platform 6000's rack-mounted keyboard/monitor. The choice of a different head node is generally related to whether more PCI slots are required for extra connectivity to external systems.

Storage

Each node must have an internal disk drive. Nodes may optionally be attached to an HP StorageWorks Modular Smart Array 1000 (MSA1000). In addition, the XC System supports high performance file I/O with the HP StorageWorks Scalable File Share (SFS). The XC System can be connected to SFS over external Ethernet connections or over the XC System Interconnect. When configuring the XC System to connect to the SFS using the XC System Interconnect, there must be enough switch ports available to accommodate the number of nodes in the XC System, as well as the number of nodes required by the SFS *(see the HP StorageWorks Scalable File Share Reference Guide)*. The switches for the combination XC/SFS system will be included as part of the Cluster Platform 3000, Cluster Platform 4000 or Cluster Platform 6000 configuration, rather than as part of the SFS.

Configuration Information

Disk Space Requirements for System Files

The storage configuration must comply, depending on whether the XC System is based on Xeon, Opteron or Itanium processor architectures, with the *Cluster Platform 3000, Cluster Platform 4000 or 6000 Reference Guides* (http://www.hp.com/techservers/clusters/ucp/index.html). Precise file system and partition size requirements are provided in the XC System Software Installation Guide, which is part of the *XC System Software Administration Documentation Kit*. Each node must have one local disk with at least 36 GB to be used for system files, swap and user data. Typical minimal disk requirements for each node are:

/	8.0 GB
/boot or /boot/efi	0.2 GB
/var	4.0 GB
/cluster	2.0 GB
swap	Same as memory size to closest 2 GB.

These sizes are approximate; actual sizes may vary depending on the user's system environment, configuration, and software options.

Memory Requirements

At least 1 GB of memory per processor is required on each node of the XC System.

Software Requirements

XC System Software is self-contained and does not generally require the installation of any other software, though most sites will want to augment it with additional compilers and tools (see the following). One exception is HP Mathematical Libraries (MLIB), which depends on specific compiler runtime libraries.

For XC6000, Itanium, MLIB is built with Intel V7.1 and V8.0 Fortran and C compilers. For XC4000, Opteron, MLIB is built with PGI V5.1 Fortran and C compliers. Note, when linking with the PGI C compiler the PGI Fortran runtime libraries may be required. For XC3000, Intel Xeon, MLIB is built with EM64T Intel Fortran and C compilers.

MLIB for XC4000 (Opteron) can be utilized on XC3000 (Intel Xeon) hardware if the PGI compilers are installed. MLIB for XC3000 (Intel Xeon) can be utilized on XC4000 (Opteron) hardware if the Intel EM64T compilers are installed.

On all XC architectures, it is recommended to link MLIB with the respective Fortran or C compilers instead of using the linker (ld) directly.

Cluster Platform 6000 systems (Integrity servers)
- Intel Itanium compilers
  - l_cc_pc_7[1].1.038.tar
    - support package ID l_cc_pc_7.1.038
    - intel-isubh7-7.1-42.ia64.rpm
    - intel-ecc7-7.1-42.ia64.rpm
    - (intel-eidb7-7.2.2-83.ia64.rpm)
  - l_fc_pc_7[1].1.040.tar
    - support package ID l_fc_pc_7.1.040
    - intel-ifc7-7.1-41.ia64.rpm
    - (intel-eidb7-7.2.2-81.ia64.rpm)
  - l_cc_pc_8[1].0.066.tar.gz
    - support package ID l_cc_pc_8.0.066
    - intel-isubh8-8.0-57.ia64.rpm
    - intel-icc8-8.0-57.ia64.rpm
    - (intel-eidb7-7.32-114.ia64.rpm)
  - l_cc_pc_8[1].0.066_pl068.1.tar.gz
    - support package ID l_cc_pc_8.0.066_pl068.1
    - Intel Fortran compilers for Cluster Platform 6000 systems (Integrity servers with Itanium 2 CPUs)
  - l_fc_pc_8[1].0.046.tar.gz
    - support package ID l_fc_pc_8.0.046
    - intel-ifort8-8.0-57.ia64.rpm
    - (intel-eidb7-7.3.2-113.ia64.rpm)
  - l_fc_pc_8[1].0.046_pl048.1.tar.gz
    - support package ID l_fc_pc_8.0.046_pl048.1

Cluster Platform 4000 and 3000 systems (ProLiant servers with Opteron or EM64T CPUs)
- Intel EM64T compilers
  - Package ID: l_fce_pc_8.1.023
  - Package ID: l_cce_pc_8.1.023

Portland Group compilers
- Version 5.1

Software

Optional Software

For access to the HP SFS Lustre file system, it is necessary to install client software that is included with HP SFS V2.0-0.

Since XC System Software is compatible with a standard Linux distribution, third party software will work correctly with XC System Software if it works correctly with that Linux distribution. Some parallel distributed software packages have been explicitly tested with the XC System Software, since parallelism typically causes complex interdependencies between separate software components layered on top of the Linux distribution. The tested software packages (unless explicitly stated, tested on all the XC-supported hardware platforms) include:

- Quadrics SHMEM, as part of QsNet^{II} user libraries, on Itanium systems connected with the Quadrics QsNet^{II} switch
- Etnus TotalView debugger Version 6.4
- Intel MKL V6.0 <interaction with HP MPI>
- Scientific Computing Associates TCP Linda 7.1
- Intel Visual Analyzer/Tracer (formally Pallas Vampir and Vampirtrace performance analyzer) on Itanium
- Gnu make (including distributed parallel make)
- MPICH 1.2.6 (generally bypasses SLURM <pointer to information>).
- PBSPro (generally bypasses SLURM <pointer to information>).

Software Licensing Information

Use of XC System Software is subject to an HP base license, an HP quantity-specified license, and the license terms in files in the physical media. Using the software indicates acceptance of all of these license terms. XC System Software contains portions of third-party software, some of which is proprietary to that third party, and some of which is open source software, including but not limited to: Linux kernel and packages, Quadrics libraries, Myricom libraries, and MPI libraries (the Bundled Products). License and warranty terms associated with the Bundled Products appear in files with the Bundled Products, which are distributed, for example, in a CD kit.

The HP Software License Terms apply to that binary-only portion of the XC System Software that is proprietary to HP or to a third-party and does not apply to or override, expressly or by implication, the terms associated with the Bundled Products. The HP Software License Terms include a license grant in which HP grants you a nonexclusive, non-transferable license to Use one copy of the XC System Software on the Quantity Of Processors contained in a Conformant System. A "Conformant System" is an HP computer or a connected collection of HP computers that conform to the hardware configuration specifications provided in the Software associated documentation. The quantity of processors for which the license applies (Quantity Of Processors) is the total of the number of processors specified as part of the HP quantity-specified license. "Use" means storing, loading, installing, executing, or displaying the Software. Other terms of the HP Software License are provided on the license agreement that is delivered with the XC System Software.

For more information about the Hewlett-Packard Company licensing terms and policies, contact your local Hewlett-Packard office.

License Management Facility Support

XC System Software supports the FlexLM License Management Facility. The FlexLM license key is provided upon presentation of a valid license key request form. This form is delivered with the software license (see the ORDERING INFORMATION section of these QuickSpecs). For more information about installing the XC System Software license keys, refer to the *XC System Software Installation Guide*.

Software and Services Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. For example, HP warrants only that the Software media will be free of physical defects for a period of ninety (90) days from delivery. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The following information is valid at time of release. Please contact your local Hewlett-Packard office for the most up-to-date information.

Combined Software License and Support Parts

On an XC System, all of the processors must be licensed. Every XC System must also include a base license for the appropriate processor architecture. The licenses can be combined. For example, a system with 17 Integrity nodes (each 2 processors) could be licensed by combining two 16-processor licenses with one 2-processor license and one Itanium base license. Alternatively, if you plan to upgrade your system later, you could order a single 64-processor license plus an Itanium base license, which would then leave some room for later system expansion without requiring additional software licenses. Licenses are combined with software support. The licenses provide permanent rights to use a specific version of the software. Standard software support includes 9x5 telephone support, media kits for new distributions, and rights to new versions for the specified period of time (one year or three years). 24x7 software support extends the telephone support window to round-the-clock access. See the section below on Software Product Services.

All of the part numbers shown below can be ordered through HP's FUSION ordering system. Part numbers with asterisks will be phased out in August 2005 and replaced with new parts shown in the table below. These new parts will be orderable through HP's VISTA ordering system as well as through HP's FUSION ordering system.

BASE LICENSE KITS		
Part Number	Description	Short Product Name
BA525A #110*	Unlimited Opteron base license, media and hardcopy docs	HP XC for OPTERON MEDIA, MANUAL, LTU
BA525A #210*	Unlimited Itanium base license, media and hardcopy docs	HP XC for ITANIUM MEDIA, MANUAL, LTU
BA525A #310*	Unlimited Xeon base license, media and hardcopy docs	HP XC for XEON MEDIA, MANUAL, LTU
LICENSES BUNDLED WITH 1 YEAR 9 X 5 SUPPORT AND UPDATES		
Part Number	Description	Short Product Name
BA525A #010*	Unlimited 1-processor use license and 1 year software support	HP XC 1 PROC LTU 1 YR SW SUPPORT
BA525A #011	Unlimited 2-processor use license and 1 year software support	HP XC 2 PROC LTU 1 YR SW SUPPORT
BA525A #012*	Unlimited 16-processor use license and 1 year software support	HP XC 16 PROC LTU 1 YR SW SUPPORT
BA525A #013	Unlimited 64-processor use license and 1 year software support	HP XC 64 PROC LTU 1 YR SW SUPPORT
BA525A #014*	Unlimited 128-processor use license and 1 year software support	HP XC 128 PROC LTU 1 YR SW SUPPORT
BA525A #015	Unlimited 256-processor use license and 1 year software support	HP XC 256 PROC LTU 1 YR SW SUPPORT
BA525A #016	Unlimited 512-processor use license and 1 year software support	HP XC 512 PROC LTU 1 YR SW SUPPORT
BA525A #017	Unlimited 1024-processor use license and 1 year software support	HP XC 1024 PROC LTU 1 YR SW SUPPORT
BA525A #018	Unlimited 2048-processor use license and 1 year software support	HP XC 2048 PROC LTU 1 YR SW SUPPORT
BA525A #019	Unlimited 4096-processor use license and 1 year software support	HP XC 4096 PROC LTU 1 YR SW SUPPORT
LICENSES BUNDLED WITH 1 YEAR 24 X 7 SUPPORT AND UPDATES		
Part Number	Description	Short Product Name
BA525A #030*	Unlimited 1-processor use license and 1 year 24x7 software support	HP XC 1 PROC LTU 1 YR 24x7 SW SUPPORT
BA525A #031	Unlimited 2-processor use license and 1 year 24x7 software support	HP XC 2 PROC LTU 1 YR 24x7 SW SUPPORT

BA525A #032*	Unlimited 16-processor use license and 1 year 24x7 software support	HP XC 16 PROC LTU 1 YR 24x7 SW SUPPORT
BA525A #033	Unlimited 64-processor use license and 1 year 24x7 software support	HP XC 64 PROC LTU 1 YR 24x7 SW SUPPORT
BA525A #034*	Unlimited 128-processor use license and 1 year 24x7 software support	HP XC 128 PROC LTU 1 YR 24x7 SW SUPPORT
BA525A #035	Unlimited 256-processor use license and 1 year 24x7 software support	HP XC 256 PROC LTU 1 YR 24x7 SW SUPPORT
BA525A #036	Unlimited 512-processor use license and 1 year 24x7 software support	HP XC 512 PROC LTU 1 YR 24x7 SW SUPPORT
BA525A #037	Unlimited 1024-processor use license and 1 year 24x7 software support	HP XC 1024 PROC LTU 1 YR 24x7 SW SUPPORT
BA525A #038	Unlimited 2048-processor use license and 1 year 24x7 software support	HP XC 2048 PROC LTU 1 YR 24x7 SW SUPPORT
BA525A #039	Unlimited 4096-processor use license and 1 year 24x7 software support	HP XC 4096 PROC LTU 1 YR 24x7 SW SUPPORT
LICENSES BUNDLED WITH 3 YEARS 9 X 5 SUPPORT AND UPDATES		
Part Number	Description	Short Product Name
BA525A #050*	Unlimited 1-processor use license and 3 years software support	HP XC 1 PROC LTU 3 YR SW SUPPORT
BA525A #051	Unlimited 2-processor use license and 3 years software support	HP XC 2 PROC LTU 3 YR SW SUPPORT
BA525A #052*	Unlimited 16-processor use license and 3 years software support	HP XC 16 PROC LTU 3 YR SW SUPPORT
BA525A #053	Unlimited 64-processor use license and 3 years software support	HP XC 64 PROC LTU 3 YR SW SUPPORT
BA525A #054*	Unlimited 128-processor use license and 3 years software support	HP XC 128 PROC LTU 3 YR SW SUPPORT
BA525A #055	Unlimited 256-processor use license and 3 years software support	HP XC 256 PROC LTU 3 YR SW SUPPORT
BA525A #056	Unlimited 512-processor use license and 3 years software support	HP XC 512 PROC LTU 3 YR SW SUPPORT
BA525A #057	Unlimited 1024-processor use license and 3 years software support	HP XC 1024 PROC LTU 3 YR SW SUPPORT
BA525A #058	Unlimited 2048-processor use license and 3 years software support	HP XC 2048 PROC LTU 3 YR SW SUPPORT
BA525A #059	Unlimited 4096-processor use license and 3 years software support	HP XC 4096 PROC LTU 3 YR SW SUPPORT
LICENSES BUNDLED WITH 3 YEARS 24 X 7 SUPPORT AND UPDATES		
Part Number	Description	Short Product Name
BA525A #070*	Unlimited 1-processor use license and 3 years 24x7 software support	HP XC 1 PROC LTU 3 YR 24x7 SW SUPPORT
BA525A #071	Unlimited 2-processor use license and 3 years 24x7 software support	HP XC 2 PROC LTU 3 YR 24x7 SW SUPPORT
BA525A #072*	Unlimited 16-processor use license and 3 years 24x7 software support	HP XC 16 PROC LTU 3 YR 24x7 SW SUPPORT
BA525A #073	Unlimited 64-processor use license and 3 years 24x7 software support	HP XC 64 PROC LTU 3 YR 24x7 SW SUPPORT
BA525A #074*	Unlimited 128-processor use license and 3 years 24x7 software support	HP XC 128 PROC LTU 3 YR 24x7 SW SUPPORT
BA525A #075	Unlimited 256-processor use license and 3 years 24x7 software support	HP XC 256 PROC LTU 3 YR 24x7 SW SUPPORT

Ordering Information

BA525A #076	Unlimited 512-processor use license and 3 years 24x7 software support	HP XC 512 PROC LTU 3 YR 24x7 SW SUPPORT
BA525A #077	Unlimited 1024-processor use license and 3 years 24x7 software support	HP XC 1024 PROC LTU 3 YR 24x7 SW SUPPORT
BA525A #078	Unlimited 2048-processor use license and 3 years 24x7 software support	HP XC 2048 PROC LTU 3 YR 24x7 SW SUPPORT
BA525A #079	Unlimited 4096-processor use license and 3 years 24x7 software support	HP XC 4096 PROC LTU 3 YR 24x7 SW SUPPORT
*REPLACEMENT PARTS (TO BE ACTIVATED IN AUGUST 2005)		
Current Part Number	New Part Number (orderable with VISTA or FUSION)	Short Product Name
BA525A #110	BA573A	HP XC for OPTERON MEDIA, MANUAL, LTU
BA525A #210	BA574A	HP XC for ITANIUM MEDIA, MANUAL,LTU
BA525A #310	BA575A	HP XC for XEON MEDIA, MANUAL, LTU
BA525A #100	BA571A	HP XC OPTERON MEDIA
BA525A #200	BA572A	HP XC ITANIUM MEDIA
BA525A #300	BA573A	HP XC XEON MEDIA
BA525A #0BE	BA576A	HP XC ADMIN MANUAL
BA525A #0BF	BA577A	HP XC USER MANUAL
BA525A #010	BA578A #100	HP XC 1 PROC LTU 1 YR SW SUPPORT
BA525A #030	BA578A #200	HP XC 1 PROC LTU 1 YR 24x7 SW SUPPORT
BA525A #050	BA578A #300	HP XC 1 PROC LTU 3 YR SW SUPPORT
BA525A #070	BA578A #400	HP XC 1 PROC LTU 3 YR 24x7 SW SUPPORT
BA525A #012	BA579A #100	HP XC 16 PROC LTU 1 YR SW SUPPORT
BA525A #032	BA579A #200	HP XC 16 PROC LTU 1 YR 24x7 SW SUPPORT
BA525A #052	BA579A #300	HP XC 16 PROC LTU 3 YR SW SUPPORT
BA525A #072	BA579A #400	HP XC 16 PROC LTU 3 YR 24x7 SW SUPPORT
BA525A #014	BA580A #100	HP XC 128 PROC LTU 1 YR SW SUPPORT
BA525A #034	BA580A #200	HP XC 128 PROC LTU 1 YR 24x7 SW SUPPORT
BA525A #054	BA580A #300	HP XC 128 PROC LTU 3 YR SW SUPPORT
BA525A #074	BA580A #400	HP XC 128 PROC LTU 3 YR 24x7 SW SUPPORT

Ordering Information

Distribution Media

XC System Software, including sources for Open Source software, is available on DVDs and its documentation is available on CD-ROM. The DVDs and CD-ROM are both included in the architecture-specific media kits, XC System Software V2.0A Media Kit for AMD Opteron™ (BA525A #100*), XC System Software V2.0 Media Kit for Intel® Itanium® (BA525A #200*) and XC System Software V2.0A Media Kit for Intel® Xeon™ (BA525A #300*). Documentation in printed format can be ordered separately; see the SOFTWARE DOCUMENTATION section of these QuickSpecs for more information. Patch kits may be required, and can be downloaded from the website http://www1.itrc.hp.com/.

Software Documentation

Documentation is included in the same Media kits as the software. The HP System Software Documentation CD-ROM includes LSF, HP MPI and HP MLIB documentation in addition to all of the material available separately in hardcopy format. Hardcopy documentation kits are available as listed below with their component manuals. Each hardcopy documentation kit includes the HP System Software Documentation CD-ROM.

XC System Software Administration Documentation Kit: BA525A #0BF*

- *HP XC System Software Installation Guide*
- *HP XC System Software Administration Guide*
- *HP XC System Software Hardware Preparation Guide*
- *Linux Administration Handbook*
- *HP XC System Software Documentation CD-ROM*

XC System Software User Documentation Kit: BA525A #0BE*

- *HP XC System Software User's Guide*
- *HP XC System Software Documentation CD-ROM*

Hardcopy documentation is also available for HP MPI and HP MLIB.

- *HP MPI User's Guide (9th Edition) : B6060-96018*
- *HP MPI Release Notes Version 2.1.1: T1919-90008*
- *HP MLIB User's Guide (6th Edition): B6061-96023*
- *HP MLIB Release Notes Version 1.5: T1918-96025*

Software Product Services

For the purposes of Hewlett-Packard Software Product Services, the product is regarded as XC System Software only if it has not been modified by the customer. Although the software includes licenses permitting modification of certain parts of the software, such modifications - unless they are provided by HP under a support agreement, or provided by HP as a licensed product - will result in a variation of the software that may no longer be called XC System Software, and that therefore is not the subject of any standard service agreement for XC System Software.

Standard software support includes customer access to technical resources during standard hours (see later), problem analysis, escalation management and resolution. HP also provides unlimited access to an electronic facility that includes a knowledge database with known symptoms and solutions, software product descriptions, specifications, and technical literature. In addition, HP will also make available certain software patches, including security patches, to the XC System Software. During the term of a standard software support contract, a customer is entitled to receive new versions of the software. With standard software support, customers can access technical resources via telephone, electronic communications or FAX where available, during standard business hours on standard business days, including the hours of 8:00 am and 5:00 pm, Monday through Friday excluding HP holidays. 24x7 software support extends the access-window to 24 hours a day, from Monday through Sunday, including holidays. Business terms and conditions governing software services can be found at the HP website http://legal.hp.com/legal/files/Standard_Agreements.asp.

Ordering Information

A variety of customer service options (BA525A #***, where the asterisks denote variant fields) are available from Hewlett-Packard for XC System Software. Service offerings are also available for additional software packages that may be distributed along with XC System Software but that are otherwise not included as part of the XC System Software. For more information, contact your local Hewlett-Packard office.

There are also Consulting and Integration Services available for XC System Software and listed below.

Product Number	Description	
U5628A 001	XC - one day on-site systems software knowledge transfer and five hours of customer integration management OR installation (factory or on-site) of XC software on 1-32 nodes	
U5628A 002	XC - two days on-site systems software knowledge transfer and six hours of customer integration management OR installation (factory or on-site) of XC software on 33-128 nodes	
U5628A 003	XC - three days on-site systems software knowledge transfer and six hours of customer integration management OR installation (factory or on-site) of XC software on 129-256 nodes	
U5628A 004	XC - five days on-site systems software knowledge transfer and six hours of customer integration management OR installation (factory or on-site) of XC software on 257-512 nodes	
U5628A 005	XC - ten days on-site systems software knowledge transfer and twelve hours of customer integration management	
Product Number	Description	Duration
U5617A	XC Implementation Program Management	40 hours implementation management consulting
U5618A	XC cluster systems software QuickStart	5 days on-site consulting and six hours management coordination
U5619A	XC Cluster Applications Migration, Development, and Optimization QuickStart	5 days on-site consulting six hours management coordination
U5626A	XC Applications Programming and Migration	2 days on-site formal customer training
U5627A	XC Cluster Systems Administration Course	4 days on-site formal customer training
U5620A	XC Implementation Program Management	80 hours implementation management consulting
U5621A	XC Cluster Systems QuickStart	10 days on-site consulting and twelve hours management coordination
U5622A	XC Cluster Applications Migration, Development, and Optimization QuickStart	10 days on-site consulting and twelve hours management coordination
U5626A	XC Applications Programming and Migration	2 days on-site formal customer training
U5627A	XC Cluster Systems Administration Course	4 days on-site formal customer training

For more details, please see your HP Consulting Services representative.

Ordering Information

© 2005 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Quadrics® and QsNet^{II}™ are trademarks or registered trademark of Quadrics, Ltd.

InfiniBand™ is a trademark the InfiniBand® Trade Association.

Myrinet® and Myricom® are registered trademarks of Myricom Inc.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries and are used under license. Pentium is a U.S. registered trademark of Intel Corporation.

Restricted Rights Legend
Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 for DOD agencies, and subparagraphs (c) (1) and (c) (2) of the Commercial Computer Software Restricted Rights clause at FAR 52.227-19 for other agencies.

HEWLETT-PACKARD COMPANY
3000 Hanover Street
Palo Alto, California 94304 U.S.A.

Use of this QuickSpecs and media is restricted to this product only. Additional copies of the programs may be made for security and back-up purposes only. Resale of the programs, in their present form or with alterations, is expressly prohibited.

AV-RWJEC-TE June, 2005